

METHOD AND APPARATUS FOR AUTOMATIC PLACEMENT OF ADVERTISING

Applicants hereby claim the benefit of priority of U.S. provisional patent application serial number 60/048940, filed 16 June 1997 and U.S. provisional patent application serial number 60/049877 filed 17 June 1997.

Field of the invention

This invention relates to computer systems for automated replacement of advertisements in scarce media. In particular, the present invention relates to the automated selection of direct advertising in advertising campaigns in order to utilize scarce media space efficiently and optimize revenue.

Background of the invention

Scarce media is defined as media providing limited amount of space. For example, newspaper and magazines have

limited amount of page space; radio and television have a limited amount of air time slots; Internet web pages have limited amount of viewable screen space and/or bandwidth; and so on. Media providers (such as magazines and newspaper publishers, radio and television broadcasters, Internet content providers and advertising networks) commonly bill advertisers based on two prevalent pricing models.

The first pricing model, based on the number of impressions delivered, is often referred to as "CPM" pricing (i.e. cost per thousand impressions). In CPM pricing, an advertiser pays a media provider a fixed amount based on the circulation of their advertisement, regardless of whether the advertising campaign is successful.

The second pricing model, based on explicit customer feedback, is sometimes referred to as "Direct Response" or "CPA" advertising (i.e. cost per actions). CPA pricing is distinguished from CPM pricing, because in CPA pricing, a marketer compensates a media provider a variable amount based on the success of the advertising campaign. Successful ad campaigns are measured, for example, by the number of survey forms filled out, leads generated, sales transacted, software downloaded, and so on.

Whereas in CPM pricing, an advertiser is entirely at risk for an unsuccessful advertising campaign, in the case of CPA pricing, the risk of an ad campaign's success is placed wholly on the media provider. It is therefore an important objective of media providers to select successful direct advertising campaigns in order to utilize scarce media space efficiently and optimize revenue.

A significant economic problem exists for media providers billing on CPA models, wherein two or more direct advertisers wish to advertise on the same scarce space and only one can be accommodated. The media provides would normally prefer to allocate its space to direct advertisers so as to maximize revenue. A revenue maximizing strategy dictates that media providers allocate space to the most successful direct advertising campaign, however, successful campaigns are not known on a pro forma basis. Consequently, media providers are currently burdened with the task of estimating the performance of direct advertising campaigns. Selection of direct advertising, which is currently done in an ad hoc fashion, is usually a manual process.

Further, current media allocations do not make reference to, or precisely measure, the "context" in which each direct advertisement performs best. Context is defined as the particular user to which the Advertisement is shown; and/or the particular date and/or time during which the Advertisement is displayed; and/or the particular "Division of Media" (a term defined below) in which the Advertisement is placed. Current selection of direct response advertising for CPA models is therefore highly inefficient and prone to error, often resulting in sub-optimal allocations of scarce media. Poorly performing direct advertisements represent a cost of opportunity and therefore an economic loss for the media provider. Conversely, better performing direct advertisements would generate more revenue and make better use of scarce media space.

In summary, there is an important need for a method and apparatus for improved selection of direct response advertisements and more efficient utilization of scarce media in direct response advertising.

Summary of the invention

The present invention provides a method and apparatus for efficient utilization of direct advertising space in scarce media. An advertising server (also referred to herein as an advertisement server) selects a direct advertisement, based on certain criteria, for automatic replacement in scarce media. Transaction results of the direct advertisement placement are reported back to the advertising server. In one embodiment, the direct advertiser's server reports transactions back to the advertising server by email. In a second embodiment, a direct proxy server mirrors the direct advertiser's server, including transaction processing. The direct proxy server reports the results of transactions back to the advertising server (and associated accounting system) without requiring an email program modification of the remote direct advertiser's web site. The use of a direct proxy for reporting transactions provides an independent audit of transactions at the remote direct advertiser's web site. In either embodiment the feedback of the results of direct advertisement transactions is used to provide an efficient utilization of direct advertising space by way of an automated computer system with a predictive model.

As used herein, the terms "advertising server" and "advertisement server" are used interchangeably to refer to

a server that selects an advertisement for display to a user. The invention is embodied in an automated computer system with a predictive model that has any and/or all of the following process functionality:

1. Estimating the performance of particular direct advertisements using a predictive model
2. Selecting a direct advertisement or group of Advertisements that are estimated to perform optimally from a pool of direct advertisements
3. Delivering the selected direct advertisements to a medium (e.g. a web page or a newspaper)
4. Monitoring a feedback loop to determine if the user responded to the direct advertisement
5. Updating the predictive model based on the user's response
6. Calculating the actual performance of a particular Direct Response Advertisement after a statistically significant sampling period of user responses

7. Replacing poorly performing advertisements with those estimated to perform better
8. Repeating steps 1 to 7 above in an iterative process so that the media provider has exploited scarce media in an optimal manner
9. Providing reports and usage patterns to the advertiser and the media provider

The present invention provides a method and apparatus for media providers to correlate particular advertisements with particular divisions of media space, so the media provider can further optimize scarce media space. Divisions of media space are, for example, individual pages of web sites; time slots for TV and radio commercials; sections of magazines and newspapers; and so on.

The present invention provides a method and apparatus for media providers to correlate particular advertisements direct with particular users and groups of users of media, so the media provider can further optimize scarce media space. Users of media are, for example, browsers of web sites; viewers of television and CATV; listeners of radio; readers of magazines and newspapers; and so on).

The present invention provides a method and apparatus for the automatic prevention of ad "burn out." Burn out is defined as the circumstance in which a viewer sees a particular advertisement, or a type of advertisement, too many times and is thereby desensitized to it, making it unlikely that he or she will respond. The same functionality for ad selection and targeting, as described in this invention disclosure, also alleviates user burn out.

It is also an object of the present invention to provide a method and apparatus for the dynamic replacement of advertisements according to context (defined above). For example, advertisements for Christmas trees have been historically shown to perform well from late-November until December 24, whereas the same advertisements perform poorly other times of the year. The same functionality for ad selection and targeting, as described in this invention disclosure, also dynamically allocates advertisements according to Context. In this example, the media allocation based upon Context of a date would cause more Christmas tree advertisements to run from late-November until Christmas Eve, and then the system will shut off such ads. While the solution in foregoing simple example would normally be

evident to a manual operator, the important feature is that in the present invention, the dynamic media allocation is automatic and does not require manual intervention. Other more subtle contexts, for which the present invention can dynamically allocate media include, but are not limited to, particular Users and Divisions of Media.

While the preferred embodiment of this invention operates on an Internet advertising network, the invention is broadly applicable across any advertising system involving scarce media and customer feedback. Applicable advertising systems include, but are not limited to, direct mail, telemarketing, television, radio, print, and CATV, etc.

Brief description of the drawings

Figure 1 is a block diagram of a system for automatic placement of direct advertisements in accordance with the present invention.

Figure 2 is a flow chart diagram of the process for processing feedback transaction data for managing the

placement of direct advertisements in accordance with the present invention.

Figure 3 is a flow chart diagram for selecting a direct advertisement based on various criteria in accordance with the present invention.

Figure 4 is block diagram of an alternate embodiment of a system for automatic placement of direct advertisements, in accordance with the present invention.

Figure 5 is a block diagram of a direct proxy server for use in conjunction with the present invention.

Detailed description of the preferred embodiments

Figure 1 is an architectural diagram of the present invention. A predictive model 10, run on a server is shown having inputs including a database of direct advertisements 14 ("Offers" in the diagram), context 16 (i.e. "Database of background Data on Users, Divisions of Media Advertisers, and other Context), Database of Historical Results of Direct Advertisements Advertisers 20, as well as information on a

particular user 22 to be served an advertisement. The predictive model 10 processes all of the foregoing inputs and outputs a single direct advertisement, or a ranked order of direct advertisements 12, to which the user is most likely to respond. The predictive model's suggestion is optionally filtered 18 by additional criteria (for example, "is the advertiser's account in good standing?" "Has the media provider blocked a particular advertisement?" etc.).

The direct advertisement 18 that is selected is dynamically delivered through a medium 26 to the user for him or her to view through a web browser. The user may then interact with the direct advertisement by means of a feedback loop 23 to a commerce engine 24 (a commerce engine is hereby defined as any electronic or physical infrastructure that facilitates a transaction). For example, Internet forms, fax-back systems, toll-free numbers, direct mail postcards, interactive television buttons, etc. are all types of commerce engines 24 in that they each enable users to respond explicitly to a direct response advertisement. Transactions are hereby defined any form of explicit feedback which a user may convey by means of a commerce engine. For example, filling out a form, placing an order, supplying a credit card number, completing

a survey, providing a survey or lead form, executing a software download, etc. are all forms of transactions.

Figure 1 also shows feedback 25 from the commerce engine 24 to the database of historical results on advertisements 20, Users, and Context. This feedback may include detailed information about a particular user's response to a direct advertisement, as well as the context, under which the response was obtained. Alternatively, the feedback loop may include a subset of the preceding information. Or the feedback loop may convey information that the user did not respond to the advertisement, if such was the case. The feedback information is used by the predictive model 10 to further refine future predictions about the optimal advertisements to deliver and maximal utilization of scarce media space.

Figure 2 illustrates the process by which a user performs a transaction with a commerce engine and through the feedback loop, reports back to the predictive model. After the user takes an action 34 and the revenue to the advertiser is calculated 38(i.e. the "payout"), the database monitoring ad delivery is automatically updated 36, 40. If the particular direct advertisement exceeds a pre-set limit

44, (for example, the maximum number of conversions for the ad, for the user, for the time period, etc.) further distribution of the advertisement is disabled 42. A conversion is a successful transaction, which depending upon the context may be any one of a sale, a click through, an order, a download, etc. The feedback process enables the predictive model, as described in more detail below, to work in conjunction with manual pre-sets on direct advertising campaigns.

The preferred embodiment of the present invention is implemented on an Internet advertising network, in which the media providers are Internet Web sites and the direct advertisers employ banner advertisements appearing for example on web pages to promote their products and services.

Both the advertisement and predictive model servers are programmed database systems running on top of Windows NT. The advertisement servers contain a database of the information necessary to generate the various banner or other types of advertisements that may appear on a web page. Each advertisement is assigned a unique address and is selected based upon the various criteria for the predictive model.

Advertising management, accounting, and monitoring are also custom-programmed systems, which are accessed through the same NT-based computer network. Database information for the predictive model, users, and context is stored on an external programmed database, which is preferably accessed through an Oracle database or other SQL database.

The predictive model is programmed and currently makes use of the following context information that may be associated for example with the Internet: user id (TCP/IP address for example), ad id (which may be arbitrarily assigned), page id (such as a World Wide Web page), site id (such as a World Wide Web site), time of day (either at the site of the computer running the predictive model or the time based on the TCP/IP address), day of the week, and date. Techniques for identifying the user ID are described in U.S. patent application No. 08/738,634 filed on October 29, 1996, pages 8, line 2 through page 9, line 10 of which are incorporated by reference. The predictive model has an extendible architecture to make use of additional context information as well.

Empirically measured data and/or initial suppositions about each advertisement are used to initially establish the predictive model to assign weights for different advertisements. For example, the time of advertisements are more favorably run may depend upon seasonal factors such as advertisements for summer clothing being more heavily favored in late spring and early summer. Similarly, the predictive model may also assign different weights to different locations where the user is located. For example, advertisements for gardening tools may be more heavily weighted when the user is viewing a page connected to gardening or home improvement and less heavily weighted towards selection when the user is viewing a page regarding for example fashions. Similarly, the time of day, the frequency of display to either all users or a given user, or any other factor such as those discussed at page 13, line 15 through page 14, line 2 of the above reference application. The predictive model may also be modified by examining the results of the CPA to maximize revenue by, for example, searching for advertisements that seem to be outperforming other advertisements in the same category of goods or services.

The advertisement servers deliver stored, targeted direct advertisements to users based on the recommendation

of the predictive model as well as additional advertisement selection information that is manually configured through the advertisement management system. The various servers for the prediction models, the advertisement servers, the commerce engines, may be in located in one physical location and connected over a LAN or they may be located in various locations and communicating over a WAN or the Internet.

Figure 3 displays the process by which the predictive model selects advertisements. In the process flow chart of figure 3, a direct advertisement is selected based on various criteria. The advertisements in this case are "offers", such as direct ads for sales of goods or services, but can also be offers to fill out survey forms, generate leads, download software, and the like. In particular, after a request for a direct advertisement offer 54 is received, user information is looked up at step 56. User information may for example, consist of the number of times the user has seen each offer.

Each of the current offers is reviewed by calculating the expected return at step 66. Some of the decisions that the predictive model bases its advertisement selection include user characteristics, past exposure to the

advertisement, historical statistical performance in the context of the user and the event; and/or the payment rate for the offer. Considering all the factors known about the user, the offer, the probability of conversion and the cost per action, the expected return for the offer is calculated. If the expected return is higher than previously rated offers at step 68, the offer is marked as the best offer at step 70. The current offer counter is incremented at step 58, and the next offer is reviewed at steps 66, 68 and 70. When all offers have been reviewed, at step 60, the best offer is sent and shown to the user at step 62.

The flow chart of figure 3 is also applicable for media content that the user requests (i.e. "pulled" content) as well as that which is broadcast to him or her (i.e. "pushed" content).

Users view Direct Advertisements as banner ads on Web Pages by using a browser such as Netscape or Internet Explorer. Direct Transactions are initiated whenever users click on these banner ads. When users click on such ads, their browsers are automatically redirected to a Commerce Server where they can complete an Action. The Commerce Server may be located at the advertiser's web site or may be

co-located at the location of the advertisement servers or the predictive model server. In response to a successful action such as a sale, an e-mail message is generated back from the advertiser's web site to the predictive model server.

Actions include, for example: completing a survey, supplying lead information, providing a credit card number, executing a software download, making a purchase, viewing a slide show, playing an online computer game, etc. The Commerce Server hosts these actions and is optionally located either on the same premises as the advertisement servers or off site at another location and/or another Internet domain. In either case, Commerce Servers communicate back to the predictive model by means of an email API. A specification of the API may be as follows:

DoubleClick Direct Response Network

Marketer Interface Specification preliminary Overview
This document specifies an interface for the feedback of transactions to DoubleClick for cost per action (CPA) advertising purchases.

When a transaction occurs that is the result of an advertisement with DoubleClick Direct, the advertiser notifies DoubleClick of the transaction via an automated email message procedure. On the ad click-through, the advertising server redirects the user to a URL specified by the advertiser such as a commerce server operated by the advertiser or a third party such as a virtual mall. The URL will end with an event string that contains information used by DoubleClick to account for transactions. For example:

<http://www.acme.com/offer.html?<event string>>

For example, the following could represent an actual click-through URL:

<http://www.acme.com/offer.html?1;3974;123;9876543;40392855;1507>

The portion after the question mark is the event string. The marketer's commerce server should remember this string and keep it associated with the user (either via a session cookie or by embedding it in the URLs used during the session).

When a transaction occurs, the commerce server must send an email notification to the operator of the predictive model server using an e-mail address such as

event@doubleclick.net. The body of the email has the following format, with arbitrary number of <event record> fields:

```
version: 1
account: <domain>
passkey: <passkey>
order: <order name>
test: <email>
<event record>
<event record>
```

An event record is a series of lines specifying details of the transaction. The event record always starts with an "event:" field. For example:

```
event: 1;3974;123;9876543;40392855;1507
action: sale
amount: 25
```

where the number "1" represents the version, the number "3974" is the advertisement identification to indicate which advertisement was used for the click through, the "number 123" is used for identifying the web site for where the banner was located; "9876543" identifies the particular page where on the web site where the banner is located, "40392855" represents the user's TCP/IP address or the identifier in a cookie for that particular user, and "1507" is a check sum used for verifying the other information;

action indicates that a "sale" took place and the amount might indicate for example 25 dollars.

Thus, all of this information will be compiled and stored on the predictive model server's database to indicate a successful action for the particular advertisement. Further, the accounting process will track the successful action and credit (either directly or indirectly) the operator of the predictive model server's accounts.

A list of fields follows:

version: means the version number for the e-mail message and is a required field; Version should generally be 1.

account: means the site's domain name (e.g., acme.com) and is a required field.

passkey: means a password which authenticates the source of the message and is provided elsewhere in the system to provide authentication so that a malicious party cannot send false messages. This is also a required field.

order: means the name of the order as shown in the ad management server and is an optional field used if the event string is unknown to correctly assign the event information to the correct order. For example, order may be used if there are multiple insertion orders with financial contract differences.

test: means an email address to send the test results and is optional. If the test address is specified the message will be validated and the results will be sent to the email address specified. In particular, this is used to verify message formats. Data contained in a test message is not processed.

event: means the event string passed from the DoubleClick ad server and is required. If the event string is unknown, the event is defined as unknown in the message format. Unknown event strings may result if the user has disabled cookies.

action: means the type of transaction such as a "sale", a "download", an "inquiry" or other categories and is required. If for a sale, the amount of the currency such as

the total amount is included as a subfield and this subfield is required.

tran-ID: means a transaction identifier string that is optional but highly recommended to include in the message to resolve any accounting discrepancies.

Each transaction may be emailed separately, or batched and sent in a single message. Batching is recommended if transactions exceed one hundred per day, otherwise, individual emails are fine. Batches for transaction should be sent at least every 24 hours, more frequently if possible.

A sample email is shown below for three events.

To: event@doubleclick.net

Subject: event

Body: version: 1

account; acme.com

event: 1;3974;123;9876543;40392855;1507 action: sale

amount: 25 tran-id: K254133

event: unknown order: Acme Test Order action: sale amount:43

tran-id: K25413

event: 1;397416;9876544;11393257;1456 action: sale amount:
230.21.tran-id: K254135

Software run on a server that may be co-located with the predictive model software compiles monthly reports for direct advertisers and web sites. The accounting software can derive the effective CPM of direct advertising campaigns by calculating this number (CPM is a function of action price multiplied by conversion rate, both of which the system keeps track of). Custom software also allows Insertion Orders to be filled out and Sites to be added to the network. The system is highly scaleable and has an architecture that will accommodate thousands of web sites and tens-of-thousands of ads.

Direct Proxy

Figure 4 shown an alternate embodiment of the present invention shown in figure 1. Figure 4 is equivalent to figure 1 except for the addition of a direct proxy 440. In particular, a system for the delivery of advertising over networks includes a user with a browser 410 (which is shown in figure 1 as user 22). The system includes at least one affiliate web site 412 (a web site as a content provider with advertising space is a specific example of a medium 26 in figure 1). Central to the system is an advertising server 414 (which includes the predictive model 10, and advertising

selection algorithms 12, 18 in figure 1). The advertising server further communicates with a database 424 (equivalent to databases 14, 16, 20 in figure 1). Also part of the overall system is one or more direct advertiser web sites 416 (represented generally as commerce engine 24 in figure 1). Email return path 418 is a specific example of feedback 25 in figure 1, while direct proxy 440 is an alternate embodiment of feedback path 25 in figure 1.

The basic operation of the systems of figures 1 and 4 provides for the selection of advertising targeted to the user. The systems shown are particularly useful when the targeted advertising is direct advertising, which requires a direct response (also known as a "conversion") by the user.

In operation, when a user, browsing on the Internet, accesses an affiliate's web site 412 which would typically include media content and advertising space 420, the user's browser 410 generates an http message to get the information for the desired Web page. The affiliate's web site 412 in response transmits one or more messages back containing the information to be displayed by the user's browser. In addition, the advertising server 414, using a local database 424 containing advertising and user data, provides

additional information comprising one or more objects such a banner advertisement to be displayed in the advertising space 420 along with the information content provided from the affiliate web site. Upon clicking through by selecting the advertising object 420 (such as a banner), the browser 410 is connected to the direct advertiser's web site 416.

Transaction results of the direct advertisement placement are reported back to the advertising server in one of two ways. In the first embodiment, the direct advertiser's server 416 reports transactions back to the advertising server by email 418. In a second embodiment, a direct proxy server 440 mirrors the direct advertiser's server 416, including transaction processing, and the direct proxy server 440 reports the results 422 of transactions back to the advertising server 414 and database 424. The direct proxy server 440 brokers the user's session (or interaction) with the direct advertiser's server 416, and reports back 422 to an accounting system (not shown) associated with the advertising server 414

The direct proxy 440 is illustrated in further detail in figure 5. Direct proxy 540 is a process that provides a mechanism for transparently monitoring activity between an

end user's Web browser 510 and an online marketer's Web site 516. The direct proxy 540 detects 522 and records 524 the successful completion of end user transactions at a remote Web site. The direct proxy 540 supports transactions including the completion of data entry forms, initiation of software downloads, and electronic sales of products or services. Because of its flexible methodology, additional types of transactions can be supported.

The direct proxy server 540 operates as follows:

1. A Web server 512 interface to the World Wide Web using the HTTP protocol, and provides a common industry standard interface such as ISAPI (Internet Server Application Programming Interface) or CGI (Common Gateway Interface). The command parser 514 is a command interpreter that breaks up received URL strings for separate action with direct proxy 540. A proxied session is started by a command received at the command parser 514. The command that begins a proxied session has associated parameters that select specific configuration information for that session. One of the associated parameters is a Uniform Resource Locator (URL) of a Web site 516 that to be proxied.

2. The direct proxy 540 fetches (from advertiser web site 516) the document specified by the supplied URL using the HyperText Transfer Protocol (HTTP) generated by the HTTP engine 532.
3. A transaction monitor 530 to look for a pre-configured data pattern in the retrieved document scans the retrieved document. The pre-configured data pattern is derived from predetermined knowledge of the type of transactions to be performed by the direct advertiser's web server 516. If the pre-configured data pattern is detected, a conversion is deemed to have been completed successfully, and the conversion (a transaction) is recorded in a database 524.
4. The retrieved document is passed through an HTML processor 528, which is a filter process that examines most of the URLs in the retrieved document. The HTML processor 528 modifies or re-assigns URLs so that remote entities are directed back to the direct proxy 540. If the URL matches a set of pre-configured patterns, that URL is modified so that the new URL now contains a direct proxy command referring back to the direct proxy's own Web server 512, with the original URL as a parameter to

the command. Thus, each URL is effectively re-addressed in such a way that when a user clicks on a hyperlink anchored by the re-addressed URL, the user's browser will be directed back to the direct proxy 540 Web server 512.

5. Steps 2 through 4 are repeated for each subsequent HTTP document request, received by the direct proxy 540, from the user's browser. A proxied session logically ends when the user clicks on a hyperlink that has not been re-addressed by the direct proxy 540 HTML processor 528.

The direct proxy 540 also contains several features designed to improve performance and transparency in the proxying process:

Web document caching. Each requested document is cached by the direct proxy 540 so that if a document has not changed since the last request, the direct proxy 540 will use the local copy in the cache, rather than re-fetching the document from the remote server.

Exclusion of binary data. The direct proxy does not re-address URLs in the HTML source that refer to binary data, such as graphic images, Java applets, ActiveX controls, and the like. Binary objects are requested directly by the

user's browser, and thus proxy bandwidth is not consumed with handling large objects (whose content is typically irrelevant to the proxy process anyway).

Cookie handling. The direct proxy will manage cookies issued by the remote server without passing these cookies back to the user's browser. This permits the cookie-based Web applications to function successfully even if the user has cookies disabled in their browser.

As indicated above, there are two alternate embodiments of the feedback path 25 in figure 1 by which transactions at a remote direct advertiser site are reported. Figure 4 illustrates both embodiments. The first embodiment is by email feedback 418, while the second embodiment is via a direct proxy 440. The use of email to provide feedback of transactions requires that the remote web site be modified to contain the email program.

Although equivalent in end results, a direct proxy 440 report has a special advantage over an email 418 report. In the case of an email report 418, the correct reporting of remote transactions is dependent on correct functioning of the email API in the remote advertiser's web site. In the case of a direct proxy report, transaction reporting is

